FULLY REVERSIBLE IMAGE ROTATION BY 1-D FILTERING

*Laurent Condat**

Institute of Biomathematics and Biometry
Helmholtz Zentrum München,
German Res. Center for Environmental Health
Neuherberg, Germany

Dimitri Van De Ville

Biomedical Imaging Group (BIG)
École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland

ABSTRACT

In this work, we propose a new image rotation algorithm. The main feature of our approach is the *symmetric reversibility*, which means that when using the same algorithm for the converse operation, then the initial data is recovered exactly. To that purpose, we decompose the lattice conversion process into three successive shear operations. The translations along the shear directions are implemented by 1-D convolutions, with new appropriate fractional delay filters. Also, the method is fast and provides high-quality resampled images.

Index Terms— Rotation, resampling, shears, interpolation, fractional delay filters

1. INTRODUCTION

Rotation amounts to resample an image from a square lattice to another rotated one. For this, typical approaches are based on reconstruction: a continuous-domain representation, e.g. a bilinear or bicubic function, is constructed, that estimates the underlying (unknown) function $f(\mathbf{x})$ by means of its samples $f(\mathbf{k})$, $\mathbf{k} \in \mathbb{Z}^2$. This function is then sampled on the target lattice, to yield the new pixel values $f(\mathbf{R}_{-\theta}\mathbf{k})$, where the rotation matrix is defined as

$$\mathbf{R}_\theta = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}. \quad (1)$$

In this paper, we propose an alternative approach that is driven by the property that the rotation operation $\mathcal{R}_\theta$ has an exact inverse operation, which is achievable with the same algorithm. That is, we pursue $\mathcal{R}_\theta \circ \mathcal{R}_{-\theta} = \mathcal{I}d, \quad \forall \theta \in \mathbb{R}$, where $\mathcal{R}_{-\theta}$ should be the converse operation of $\mathcal{R}_\theta$; that is $\mathcal{R}_{-\theta} = \mathcal{S}\mathcal{R}_\theta\mathcal{S}$, where $\mathcal{S}$ is an axial symmetry through the origin. We say a rotation operator satisfying these requirement to be "symmetrically reversible". Formally, these condition is equivalent for $\mathcal{R}_\theta$ to be *orthogonal*: $\mathcal{R}_{-\theta} = \mathcal{R}_\theta^*$, its adjoint operator, is also its inverse.

*The first author is supported by the Marie Curie Excellence Team Grant MEXT-CT-2004-013477, Acronym MAMEBIA, funded by the European Commission. Contact: laurent.condat@gsf.de.

The motivation for this work was to find a method for rotating an image without losing any information, with the guarantee that the same operation applied in opposite sense will exactly recover the initial image. The orthogonality of the method may also be very useful in some applications; for instance, it becomes equivalent to apply some denoising method on the image or its rotated version.

Our method is based on two fundamental ingredients: (1) One can turn the square lattice into its rotated version in three successive shear operations; This idea, known in the literature, makes the method separable, hence, simple to implement and fast through 1-D operations along rows and columns of the image. We recall this method in the Section 2. (2) The 1-D translation operators required to implement the shears can be carried out so as to ensure the "symmetric reversibility". We propose a new specific choice of filters to this purpose, developed in Section 3. In this family, filters of arbitrary high order can be designed to guarantee high quality results, as shown by experiments in Section 5.

2. ROTATION BY THREE SHEARS

In the following, $\mathbf{x} = [x_1\ x_2]^{\mathrm{T}}$ denotes a vector of $\mathbb{R}^2$. A continuous shear corresponds to a displacement of a point $\mathbf{x}$ in a direction $\mathbf{a}$, with amplitude proportional to $\langle \mathbf{x}, \mathbf{a}^\perp \rangle$ where $\mathbf{a}^\perp = [-a_2\ a_1]^{\mathrm{T}}$. So, a shear is characterized by a matrix of the form $\mathbf{S} = \mathbf{I} + \lambda \mathbf{a}\mathbf{a}^{\perp\mathrm{T}}$, for some $\lambda \in \mathbb{R}$. It is known that the rotation matrix can be factorized in three shear matrices (see [1] and references therein):

$$\mathbf{R}_{-\theta} = \begin{bmatrix} 1 & -\tan(\frac{\theta}{2}) \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \sin(\theta) & 1 \end{bmatrix} \begin{bmatrix} 1 & -\tan(\frac{\theta}{2}) \\ 0 & 1 \end{bmatrix}. \quad (2)$$

This decomposition provides us with the following algorithm, that only involves 1-D operations:

- On each row of the image s with index k_2, perform a translation of magnitude $\tan(\theta/2).k_2$ (in the direction $\uparrow$). This amounts to estimating the sample values $f(k_1 - \tan(\theta/2).k_2, k_2)$ from the available pixel values $f(k_1, k_2)$.

- On each column of the image with index k_1, perform a translation of magnitude $\sin(\theta).k_1$ (in the direction $\rightarrow$).
- On each row of the image with index k_2, perform a translation of magnitude $\tan(\theta/2).k_2$ (in the direction $\uparrow$). The pixel values of the final image approximate the desired samples $f(\mathbf{R}_{-\theta}\mathbf{k})$.

We remark that only rotations with angle $\theta \in [-\pi/4, \pi/4]$ have to be processed by this decomposition, since the other rotations can be decomposed in such a rotation and a rotation with angle $-\pi/2$, $\pi/2$ or π, which is trivially performed.

Decomposing the rotation in three shears provides algorithms—for the same quality—being computationally less complex than their 2-D counterparts [1]. However, the method proposed in [1], implementing the 1-D translations using spline interpolation, is not reversible. For particular angles, 2-D nearest neighbor interpolation is actually symmetrically reversible, and can be computed efficiently [2]. The quality of this method is very poor, however. It should also be stressed that other decompositions of rotations have been proposed in the literature [3, 4, 5], but they use scaling operations in addition to shears, that necessarily imply some loss of information; hence, they can not be made reversible.

3. 1-D FRACTIONAL DELAY OPERATORS FOR IMAGE TRANSLATION

We now concentrate on the way to perform the 1-D translations. A translation (a.k.a. *shift* or *delay*) operator $\mathcal{T}_\tau : s \mapsto s' = s * h_\tau$ is implemented by a discrete convolution with a *fractional delay filter* h_τ. These filters have a long history in signal processing and communication systems, see [6] and references therein. So, we have to design a family of filters h_τ, for every $\tau \in \mathbb{R}$. Further on, we denote the $\mathcal{Z}$-transform of a filter h by $H(z) = \sum_{k\in\mathbb{Z}} h[k]z^{-k}$.

The rotation is symmetrically reversible (orthogonal) if and only if the translation is orthogonal, too. This is equivalent to the two conditions $1/H_\tau(z) = H_{-\tau}(z) = H_\tau(z^{-1})$: the inverse of $\mathcal{T}_\tau$ is the translation in the opposite direction, corresponding to the same operation but in reverse order on the data. This amounts for h_τ to be an *all-pass filter*:

$$|H_\tau(e^{j\omega})| = 1, \ \forall \omega \in \mathbb{R}. \tag{3}$$

So, h_τ is entirely characterized by its phase response $\theta_h(\omega)$ such that $H(e^{j\omega}) = e^{j\theta_h(\omega)}$. The phase delay $-\theta_h(\omega)/\omega$ and group delay $-\mathrm{d}\theta(\omega)/\mathrm{d}\omega$ are classical measures of quality for delay filters. Ideally, they should be constant and equal to τ in the range $\omega \in (-\pi, \pi)$.

For integer values of τ, the translation is exact with $H_\tau(z) = z^{-\tau}$. Using compositions with such shifts and the symmetry, we only have to design the filters h_τ for $\tau \in [0, 1/2]$. Then, for $\tau \in \mathbb{R}$, we select $H_\tau(z) = z^{-d}H_{|\tau-d|}(z^{\mathrm{sgn}(\tau-d)})$, where $d = \mathrm{sgn}(\tau)\lceil|\tau| + \frac{1}{2}\rceil - 1$.

We note that two classical families of filters are all-pass and may be used in our context: nearest neighbor interpolation ($H_\tau(z) = 1 \quad \forall\tau$), which is very basic and does not provide a good quality; and the sinc interpolation kernel ($H_\tau(e^{j\omega}) = e^{-j\omega\tau}$), prone to the introduction of unwanted oscillations (ringing). Moreover, the infinitely long response of sinc interpolation requires much computation time, since an implementation in the Fourier domain using FFTs is required. The filters we propose in the following are short and realizable in the spatial domain.

The only known class of realizable all-pass filters having explicit formulas of their coefficients as a function of τ is the class of Thiran filters [6], designed to have maximally flat group delay at $\omega = 0$. This asymptotic constraint privileges the accurate translation of the low frequency part of the signal. This design is in fact particularly relevant for images, since it is well known that images have their energy content essentially localized around the origin in the frequency domain. Such "maxflat" filters are also Neville filters [7]; that is, they perfectly translate polynomial signals up to some degree. The Thiran filter of order N is causal, and stable if and only if $\tau > N - 1$. For our purpose, we want to design delay filters for $\tau \in [0, 1/2]$, which does not suit this stability interval. Moreover, the constraint for the filter to be causal is not necessary for treatments on images. Thus, we design new non-causal all-pass filters having maximally flat group delay, optimized for $\tau \in [0, 1/2]$: we define, for every integer $N \geq 1$, the new filter h_τ^N by

$$H_\tau^N(z) = \frac{1 + b_1 z^{-1} + \cdots + b_{N-1}z^{-N+1} + b_N z^{-N}}{b_N z^N + b_{N-1}z^{N-1} + \cdots + b_1 z^1 + 1}, \tag{4}$$

with

$$b_k = (-1)^k \binom{N}{k} \prod_{n=0}^{N} \frac{\tau - n}{\tau - n - k} \quad \forall k \in 1..N. \tag{5}$$

h_τ^N turns out to be the Thiran filter of order N whose denominator would have been shifted to map the interval $\tau \in [N - 1/2, N]$ into the interval $\tau \in [0, 1/2]$.

Filtering a signal s of finite length T with h_τ^N is easily achieved in a single in-place backward pass:

```
for i from T-1 down to 0 do
  for k from 1 to N do
    s[i]+=b_k*(s[i-k]-s[i+k]);
```

For the translation to be reversible without expanding the size of the signal, periodic boundary conditions have to be used. The initialization of the recursive loop is done classically, using a few terms of the signals so that the error is below some prescribed level of precision. Also, extending the size of the image before rotation is necessary in order to avoid the shuffling of some parts of the images, caused by these periodic conditions. The image can be cropped back to its initial size afterwards, but the reversibility is lost in that case.

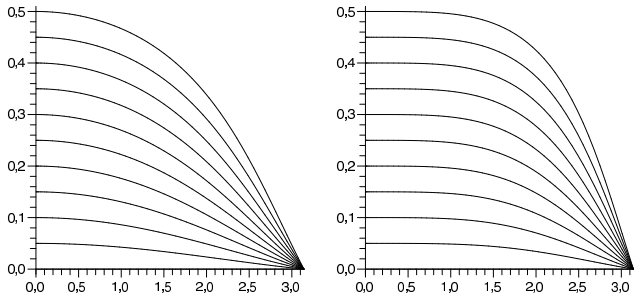

Fig. 1. Phase delay of h_τ^1 (left) and h_τ^2 (right), for several values of τ in [0,0.5].

Finally, we note that for $N = 0$ and $N \to \infty$, we obtain the translation by nearest neighbor and sinc interpolation, respectively. So, we denote these two methods by h_τ^0 and h_τ^∞ in the following. In fact, as N increases, the phase delay of the filters h_τ^N approaches more and more the ideal phase delay, as can be seen in Fig. 1.

4. IN-DEPTH ANALYSIS

First, it is worth mentionning that there is no underlying continuous model fitted on the image, as is the case with interpolation methods. The proposed approach is entirely discrete. Also, our method does not introduce any blur, since the use of all-pass filters leaves unchanged the magnitude of all frequency components along the shear directions.

An analysis of the behavior of the rotation process can be carried out in the Fourier domain, see [8, 9]. In fact, when the rotation is performed using three shears with sinc interpolation for the translations, the central part of the frequency plane is perfectly rotated, but other parts of it are displaced at other locations. This means that bandlimited images with frequency content in this central area are perfectly rotated. When using one of the proposed filters—without ideal phase response—additional (and much harder to quantify) distortions of the frequency content occur.

Concerning the computational complexity, the proposed method with h_τ^N requires only $3N$ multiplications and $6N+3$ additions per pixel. Even the simplest bilinear interpolation requires 3 multiplications and 8 additions per pixel. Moreover, the conversion process can be performed in-place on the initial data, making the use of auxiliary memory unnecessary.

5. EXPERIMENTAL VALIDATION

To illustrate the quality of our approach, we performed 9 successive rotations of angle $2\pi/9$ on natural 512×512 images. The cumulative effect was then observed by comparing the final image against the initial one. They were first extended within a larger support, before applying the rotations, and cropped to the initial size afterwards. The results are reported in Tab. 1 and are illustrated in Fig. 2.

Performing the translations using nearest neighbor interpolation clearly provides a poor quality. Our method significantly outperforms the classical bilinear interpolation, which provides blurred images. The bicubic interpolation [10] is also outperformed, using h_τ^N with $N \geq 2$. A higher order ($N \geq 5$) would be required to outperform the reference cubic spline interpolation (denoted SP3). Our method does not introduce any blur and the parameter N controls the tradeoff aliasing/ringing: the sinc filter h_τ^∞ creates ringing artifacts that spread over the entire image, while strong jittering effects appear with h_τ^0. Since these artifacts may be considered more disturbing than blur, the practitioners not interested by reversibility, who want the best tradeoff between quality and speed, may use the decomposition in shears in combination with 1-D spline interpolation, as proposed in [1]. We report the results corresponding to this approach, for cubic spline interpolation, in the fifth column of Tab. 1.

Also, we indicate for each method the computation time of one rotation, for C-code run on an Apple Mac Dual 2.7 Ghz PowerPC G5. All of our filters give computation times significantly reduced when compared with interpolation methods.

6. CONCLUSION

We proposed an image rotation method that is first and foremost symmetrically reversible (orthogonal). Thus, no loss of information is introduced during the rotation. It combines a decomposition in three shears with new all-pass fractional delay filters. As shown by practical results, our approach offers a very good tradeoff between quality and computation speed. Our approach can be extended in higher dimensions, e.g. using the decomposition of 3-D rotations in 4 shears [11].

7. REFERENCES

[1] M. Unser, P. Thévenaz, and L. Yaroslavsky, "Convolution-based interpolation for high-quality rotation of images," *IEEE Trans. Image Processing*, vol. 4, no. 10, pp. 1371–1381, Oct. 1995.

[2] S. Srinivasan, "Modulo transforms—An alternative to lifting," *IEEE Trans. Signal Processing*, vol. 54, no. 5, pp. 1864–1874, May 2006.

[3] B. Chen and A. Kaufman, "Two-pass image and volume rotation," *IEEE Workshop on Volume Graphics*, 2001.

[4] D. E. Friedman, "Two-dimensional resampling of line scan imagery by one-dimensional processing," *Photogrammetric Eng. Remote Sensing*, vol. 47, pp. 1459–1467, Oct. 1981.

[5] E. Catmul and A. R. Smith, "3-D transformation of images in scan-line order," *Proc. of SIGGRAPH*, vol. 14, pp. 279–285, 1980.

[6] T. I. Laakso, V. Välimäki, M. Karjalainen, and U. K. Laine, "Splitting the unit delay—Tools for fractional delay filter design," *IEEE Signal Processing Mag.*, vol. 13, no. 1, Jan. 1996.

	interpolation			3 shears	conversion in three shears with proposed filters				
image	bilinear	bicubic	SP3	+SP3 [1]	h_τ^0	h_τ^1	h_τ^2	h_τ^3	h_τ^∞
lena	31.57	37.13	40.26	38.94	24.74	34.74	36.90	37.89	40.95
barbara	24.67	27.55	31.29	29.32	20.30	25.11	27.67	29.36	36.67
baboon	23.13	26.46	29.16	27.61	18.86	23.92	25.69	26.61	30.01
lighthouse	24.49	29.29	33.20	31.35	19.40	27.11	29.76	31.23	36.31
goldhill	29.97	34.01	36.89	35.47	24.38	31.57	33.49	34.54	38.94
boat	28.17	32.80	35.44	34.24	22.26	30.26	32.03	32.87	35.73
time	0.28	0.63	0.65	0.10	0.05	0.08	0.09	0.10	1.17

Table 1. PSNR Results for rotations experiments on several classical images. 9 rotations of angle $2\pi/9$ are performed.

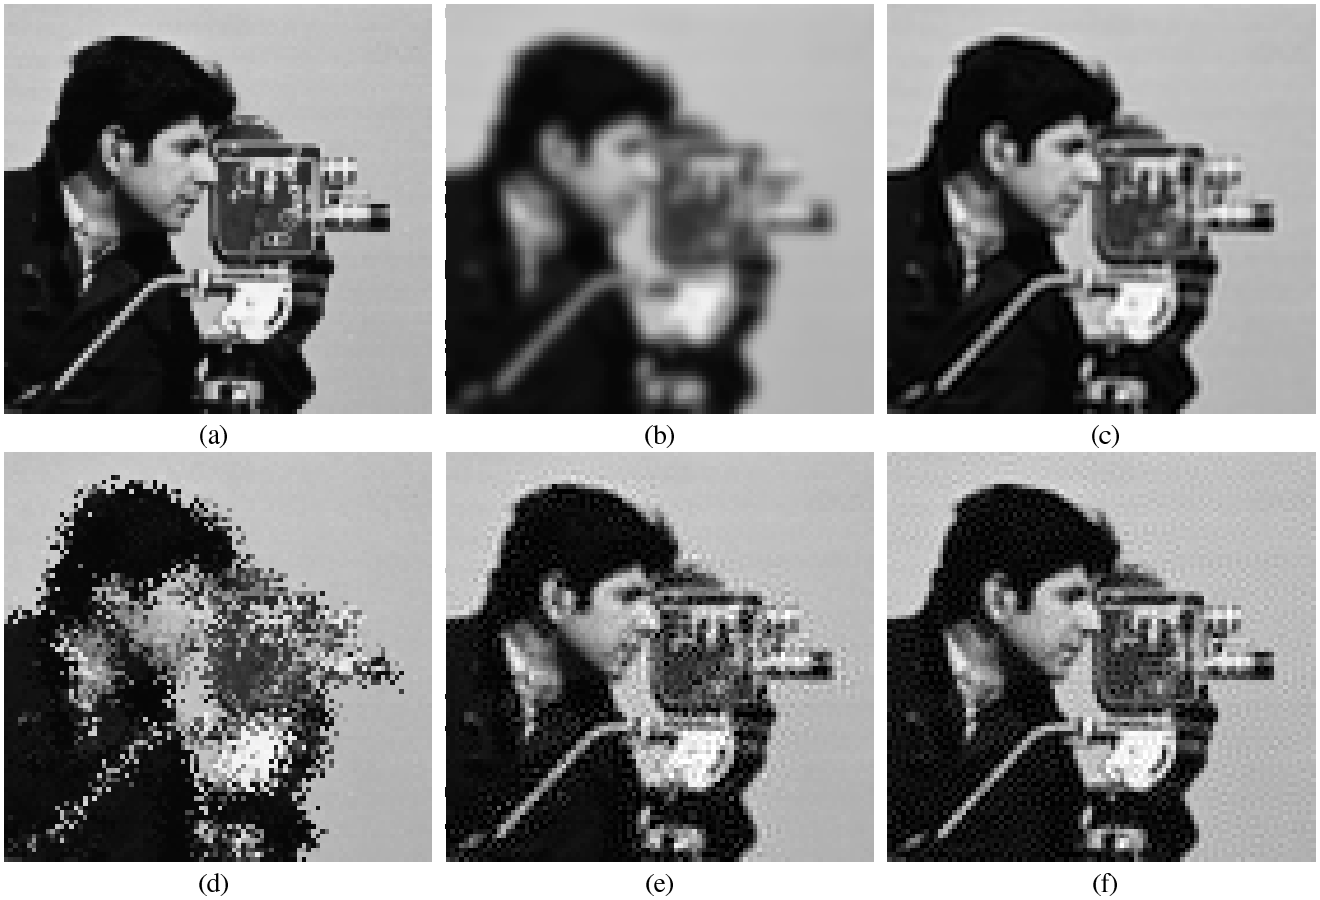

Fig. 2. Images obtained after 9 rotations of angle $2\pi/9$ on a part of the *camera* image depicted in (a), using (b) bilinear and (c) cubic spline interpolation, and the decomposition in three shears with filters (d) h_τ^0 (nearest neighbor), (e) our filter h_τ^2, (f) h_τ^∞ (sinc). The PSNRs are from (b) to (f): 21.19, 28.00, 15.63, 25.05, 28.67 dB, respectively.

[7] J. Kovačević and W. Sweldens, "Wavelet families of increasing order in arbitrary dimensions," *IEEE Trans. Image Processing*, vol. 9, no. 3, pp. 480–496, Mar. 2000.

[8] C. B. Owen, "High quality alias free image rotation," in *Proc. of 30th Asilomar Conf. on Signals, Systems, and Computers*, Nov. 1996.

[9] L. Condat and D. Van De Ville, "Reversible, fast, and high-quality grid conversions," *IEEE Trans. Image Processing*, vol. 17, no. 5, pp. 679–693, May 2008.

[10] R. G. Keys, "Cubic convolution interpolation for digital image processing," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 29, pp. 1153–1160, 1981.

[11] J. S. Welling, W. F. Eddy, and T. K. Young, "Rotation of 3D volumes by Fourier-interpolated shears," *Graphical Models*, vol. 68, no. 4, pp. 365–370, July 2006.